

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

OF

CHRISTOPHER J. MANNING

FOR

FOURIER-TRANSFORM SPECTROMETERS

Prepared by
Christopher Manning, President
Manning Applied Technology
419 South Main Street / PO Box 265
Troy, ID 83871

PROSECUTION PRACTICE

Related Applications

This application claims the benefit of Provisional Applications Serial Nos. 60/228800, filed 08/02/00, 60/242232, filed 10/17/2000 and 60/288,273 filed 05/02/2001. This application is related to pending formal applications Serial No. 09/299,022 and 09/433,964.

Previous filings by the author are included by reference for the entirety of their disclosures. The first is "Tilt-Compensated Interferometers," filed 4/26/1999, serial number 09/299,022. More are provisional applications serial number 60/107,060, filed 11/04/98, titled "FT-IR Signal Processing: Part I," serial number 60/119,429, filed 02/09/99, titled "FT-IR Signal Processing: Part II," a formal application entitled "Signal Processing for Interferometric Spectrometry" serial number 09/433,964 filed 11/04/99. Further provisional applications which are included for the entirety of their disclosures are titled "Interferometers and Interferometry," serial number 60/228800, filed 08/02/00, titled "Interferometers and Interferometry: Part 2," serial number 60/242232, filed 10/17/2000, and titled "Interferometers and Interferometry: Part 3," serial number 60/288,273 filed 05/02/2001. The book by Griffiths and deHaseth, "Fourier transform spectrometry," ISBN 0-471-09902-3, is also included for the entirety of its content.

Portions of the inventions disclosed here have been made under contract with the United States Federal Government through the Department of Defense under contracts DAAD13-01-P-0012 and DAAD-05-96-P-3102. The Government has certain rights in these inventions.

Background and Summary of the Invention

It is an object of the present inventions to provide new interferometers, which are better than prior art in respect to stability, scan speed and cost of manufacture. It is an object of the present inventions to provide signal-processing techniques for interferometry, particularly very rapid scan interferometry. It is an object of the present inventions to improve the state-of-the-art in photometric accuracy of interferometric measurements.

Brief Description of the Drawings

Figure 1 shows a diagram of a tilt-compensated interferometric spectrometer.

Figure 2 shows a diagram of the reflector/beamsplitter assembly of Figure 1.

Figure 3 shows a diagram of a Michelson interferometer incorporating refractive path difference scanners.

Figure 4 shows a diagram of a modified Michelson interferometer design incorporating refractive path difference scanners.

Figure 5 shows a diagram of a bearing and support mechanism for producing a nutating motion of a refractive scanning element.

Figure 6 shows a diagram of a remapping function applied to modify the apparent wavenumber spacing of the x-axis to actually wavenumber spacing.

Figure 7 shows a diagram of a signal processing chain that has the effect of recording the inverse transfer function of a detector signal chain.

Figure 8 shows a diagram of signals which are produced during a real-time measurement of detector saturation and correction for saturation.

Figure 9 shows a diagram of an arrangement which may be used to measure and optionally stabilize the wavelength of a tunable solid-state laser to render it an effective wavelength standard for interferometry.

Figure 10 shows a diagram of signal produced during measurement of the wavelength of a tunable solid state laser.

Figure 11 shows a diagram of a permanently-aligned, tilt-compensated four-port interferometer.

Figure 12 shows a diagram of a permanently-aligned, tilt-compensated four-port interferometer with both inputs derived from a common source, and the outputs recombined to effect optical subtraction.

Figure 13 shows an alternate geometry to that of Figure 12 in which the functions of the separate beamsplitters and combiners are done by a single assembly.

Figure 14 shows a diagram of a four-port interferometer according to the approach of Figure 13.

Figure 15 shows a diagram of an optical path difference scanner that uses gas pressure to effect a variable optical delay.

Figure 16 shows a diagram of an optical path difference scanner that uses an airbearing assembly to avoid shear alignment errors.

Detailed Description

Beamsplitter-roof reflector assembly

Larsson has described in US Patent 5,650,848 a method for tilt-compensating the scanning mechanism of an interferometer. This prior art extends ideas described by Steel in the early 1960's (see for example, "Interferometry," Cambridge University Press, 1967 W. H. Steel). A new approach substantially increases the throughput that is possible, while making it straightforward to produce very compact reflector assemblies and interferometers. The crux of the new invention is an assembly that consists of a beamsplitter rigidly mounted to a roof reflector. By roof reflector, what is meant is a pair of mirror facets which are mutually abutted and mutually perpendicular to interferometric tolerances.

A slight disadvantage of the Michelson interferometer is that the two end mirrors in the arms are susceptible to misalignment with each other and with the beamsplitter. The Jamin (1856), Solomon (US Patent 5,675,412), Turner and Mould (US Patent 5,808,739), Woodruff (US Patent 4,391,525) and several related designs provide an interferometer framework in which slight misalignments of the optical components are completely compensated. Thus, by adding one or more refractive scanning elements to such designs, it is possible to construct an interferometer that

preserves intrinsic tilt-compensation and allows variation of path difference with minimum reduction of throughput. If the refractive scanning elements have faces that are parallel and flat, then they will usually not compromise the intrinsic tilt-compensation of an interferometer. The exception is that for interferometers that do not utilize path reversal, the shift caused by such optical scanning elements can cause shear, but not angular, misalignment of the beams in the two arms.

Figure 1 shows a diagram of an interferometer which is permanently aligned and for which scanning of the optical path does not require interferometrically accurate rectilinear motion. Radiation from a source impinges on a beamsplitter coating which is supported on one side of a substrate. The beam reflected from the coating traverses to a hollow cube corner or other retroreflector where it is directed back through a portion of the beamsplitter substrate which is transparent. This traversal of the beamsplitter substrate provides compensation for the dispersion of the substrate material. The beam then propagates to a flat reflector which returns it on its path back to the beamsplitter. The radiation which passes through the beamsplitting coating traverses to a block which is rigidly affixed to the beamsplitter substrate. Preferably the block is a hollow roof reflector (two facets at 90 degrees to each other, both of these facets being oriented at 90 degrees relative to the beamsplitting coating). If these two faces are arranged to be mutually perpendicular to the beamsplitting coating of the substrate, then three facets will form a cube corner, or lateral transfer retroreflector. The relevant surfaces of the block and can be coated with reflective materials by a variety of methods. The block is understood to be of interferometrically precise construction such that after a reflection from the block the beam propagating to the flat reflector will have a wavevector that is exactly parallel to the beam from the retroreflector. Thus, the flat mirror will return the second beam back to the beamsplitter. It may be discerned that after recombination, both beams will have traversed an exactly equal thickness of the beamsplitter substrate.

Two novel methods for fabricating the roof reflector and end mirror of the interferometer are described. These are wire electron-discharge machining (EDM) and fired ceramic. It is thought that by the choice of metals which have very uniform coefficients of expansion, together with wire EDM, that excellent precision substrates can be prepared. Preferred metals are aluminum, invar, and stainless steel. EDM uses an electric arc to erode metal from a part. The electric arc produces very little damage to the part and particularly does not produce stresses in the parts that can cause adverse thermal distortion. Wire EDM can produce surface finishes that are excellent for replication. Replication is a known method for producing high-quality optical surfaces. Another preferred approach to preparation of substrates for the roof reflector is slipcasting of ceramic mixtures and molding of ceramic pastes (i.e., clays). By judicious choice of components, the thermal expansion properties of the ceramic can be tailored to be uniform in all directions and very small. In particular, the addition of ceramic microspheres is preferred because the resulting masses have uniform properties. Some ceramic compositions have particles that have non-uniform properties. If the particles are shaped such that they align preferentially (e.g., flakes or rods), then the macroscopic properties of the resulting masses may not be uniform. After casting or forming, it is understood that the resulting ceramic pieces are fired, then used as substrates for replication.

This inexpensive interferometer is suitable for use with the nutating disk refractive scanning elements described below as well as a variety of other scanning methods. The beamsplitter with anti-parallel reflector assembly may be used as a component of the optical subtraction interferometers described below. The refractive scanning elements described below be inserted at any point in the optical paths in either arm of the interferometer. Motion of the cube corner

provides a method for fine-tuning the zero path difference point of the interferometer, or it may be used as the optical path difference scanning element. Alternatively, rotating disks described in US Patent 5,898,495 may be interposed in the path between the beamsplitter and the cube corner reflector. It may be necessary to manufacture the final flat reflector in a stair step configuration to compensate the differences in path length of the two arms of the interferometer, thereby insuring that zero path difference is accessible by the scanning element or elements. Two views of the beamsplitter with roof reflector are also shown in Figure 2. These clarify the layout of the components on the beamsplitter.

Nutating disk, refractive scanning interferometers

The use of a nutating disk of transparent material as a refractive scanning element of an interferometer is disclosed herein. The principle of optical path difference scanning is that the disk is mounted such that a beam passes through it. During rotation, nutation occurs so that the angle of the disk with respect to the beam varies; consequently the optical path is varied. Such a disk must generally be supported by the edges so as to not block radiation from being transmitted through the disk. The disk is rotated approximately about the axis near where the center of the beam passes through it. The axis of rotation is not exactly perpendicular to the plane passing through the edge of the disk. Thus, the disk nutates or precesses as it rotates thereby causing the optical path through the disk to vary. The disk should have faces that are very flat and parallel such that after passing through the crystal the angle of propagation of the beam is unchanged. In cases where the beam is inverted and passed through the disk a second time, it may be acceptable for the faces to not be parallel, but they should still be flat to interferometric tolerances.

A series of methods for mounting and spinning these disks are disclosed. Two preferred materials for construction of the disks are sodium chloride and potassium bromide, hereafter NaCl and KBr. The disadvantage of these materials is that they are relatively soft. However, calculations show that the deformation of KBr even when spun at 500 revolutions per second, suitable for measuring 1000 interferograms per second, will be quite small. Other materials are also contemplated for use as nutating refractive scanning elements. In particular, germanium is good because of its stiffness and high refractive index. Because of the high refractive index of germanium it should be used with antireflection coatings. Many other materials including silica (quartz), alumina, silicon, zinc sulfide and zinc selenide may be used depending on the desired spectral range, spin rate, retardation and cost range.

The advantage of the nutating arrangement relative to other arrangements described in the literature (e.g., US Patent 5,898,495) is that the throughput can be much higher. In the case of Dybwad's design (US Patent 4,654,530) the substantial advantage of the nutating disk approach is improved duty cycle efficiency; however, the throughput is generally not significantly greater than that obtained by Dybwad. The throughput is significantly greater than that obtained by Davis (US Patent 5,883,713). The disadvantage of Davis' design is that the beam has to pass twice through a salt crystal, which is generally cubic. The duty cycle efficiency is consequently poor, as is the throughput. A point should be made that the rotation axis of the scanning element in Dybwad and Davis are both perpendicular to the scan axis described herein. In both the present case and US patent 5,898,495, the fundamental scanning mechanism is nutation. Nutation of a transparent disk allows much higher throughput than nutation of a disk mirror. The reason for the difference is that with the disk mirror the beam must be reflected back to the disk mirror for a second reflection. The length of the return path is on the order of 3 d, where d is the beam diameter. In a nutating disk design, the total length in either arm of an interferometer must be on the order of 6 to 10 d. By

using a nutating transparent disk, the total length in each arm of an interferometer can be reduced to the order of 3 d. Throughput can thus be approximately 10 times higher because it scales as the square of the number of beam diameters of path traversed in each arm.

Collins, et. al have disclosed in US Patent 3,955,083 the use of a nutating disk as an interferometric shaft angle encoder. The device differs considerably from the present disclosure, both in purpose and implementation. Collins described a disk mounted by a shaft through the center. Nutation of an optical element for the purpose of interferometric modulation has also been published by Crane in US Patent 4,743,114. Crane's art differs considerably from the present disclosure for several reasons. In particular, the motion of the beam produced by nutation of the optical element was not cancelled by a second pass in the reverse direction, but rather exploited as a means of scanning the field of view. Further, Crane's element was a Fabry-Perot etalon rather than a disk for refractive scanning of optical path difference.

A transparent nutating disk element can be used with a variety of basic interferometer designs, including the standard Michelson. Figure 3 shows the preferred embodiment with a Michelson interferometer in which one disk is located in each arm. The disks are oriented at approximately 45 degrees relative to the beams. The disks are held by their edges in strong support cylinders, which may lend some precompression to enhance the stability of the crystal in a centrifugal field. The cylinders may be mounted in ball bearings for rapid rotation. The disk is oriented at a small angle relative to the rotation such that it appears to wobble or nutate during rotation. This has the effect of altering the path length for radiation passing through the crystal. A preferred embodiment has the cylinders supported by magnetic levitation. The same means can also drive very rapid rotation, by altering the magnetic fields at the support coils in synchrony. Either one or both disk elements can be rotated. If both are rotated, they may be rotated at the same rate, or at different rates. If one is rotated at twice the rate of the other, then interleaved scans of high and low resolution will be obtained.

Figure 4 shows a diagram of a folded Michelson interferometer using a variation of an arrangement published by Steel many years ago. The advantage of passing the beam twice through the nutating disk is that any variation of angle caused by a wedge angle between the two faces of the nutating disk (i.e., non-parallelism) will be cancelled by the double pass. Thus, the nutating disk may be of less exacting manufacture. The disadvantage is that the extra length of the interferometer arms generally decreases throughput. Another disadvantage is that the disk must be twice as large to accommodate two beam diameters on its surface.

If the angle between the nutating disk and the beam is approximately 45 degrees then the variation of path difference with angle is nearly linear for small nutation or rotation angles. It is found that for a 6 mm thick KBr window having a refractive index of approximately 1.5, that the variation of path difference with angle for one disk with the beam passing in both directions is about 0.15 mm per degree of nutation (because the effect of 1 degree of nutation is -1 to +1 degree of tilt). If a disk is mounted in each arm, then the maximum optical path difference (and resolution) may be doubled.

If one disk is rotated at three times the rate of the other and is made appropriately thinner, then it may be used to linearize the change of optical path difference with respect to time. This idea may be furthered by adding two more disks and rotating them at 5 and 7 times the base rate to further linearize the change of optical path difference with respect to time. These disks would be

correspondingly thinner. It is thought that the benefits of linearizing the scan are modest, given the signal processing approaches described herein and elsewhere.

It is understood that there will be dispersion in the scanning element or elements. If a refractive scanning element is used in only one arm of the interferometer, then the interference signal may be strongly chirped. Mattson has published a method (US Patent 5,491,551) for signal processing which accounts for a strongly chirped signal.

Figure 5 shows two methods for supporting and spinning an assembly that includes a nutating disk mirror. The figure shows the housing surrounded by a ball bearing. It is understood that the rotating assembly should be balanced by drilling material from the housing so that during rapid rotation vibrational forces are minimized. The disk should be housed in a suitable material. Carbon fiber composites are preferred because of their lightweight, high stiffness and small coefficient of thermal expansion. Many other materials may be used for the same purpose. Titanium is thought to be a good choice because of its excellent mechanical properties and lightweight.

Signal processing for Dispersion

A pending application by the author describes methods of signal processing which are also appropriate for use with refractively scanned interferometers. The general approach favored by the author is to record both the laser and infrared channels separately and simultaneously using suitable analog-to-digital converters or other technology such as that proposed by Mertz circa 1965. These ADC's will be clocked from a common source so that they produce a set of matched simultaneous samples. An inverse transfer function may be applied to both channels to provide delay matching. The laser signal is then processed, preferably with a Hilbert transform or filter, to extract an accurate estimate of phase for each data point. The mirror position is proportional to this phase value. This set of data can be used together with the corresponding samples of the infrared channel to do either an unevenly-spaced Fourier transform (i.e., Lomb periodogram), or a spline or polynomial fit with subsequent conventional Fourier transform. Brault's method may also be used to produce excellent results. If the interferometer does not exhibit dispersion, the resulting spectral data are equivalent to those produced by conventional signal processing means. The signal processing steps detailed below will remove the effects of dispersion.

In the case of a refractively scanned interferometer, spectral distortion will result from dispersion. All refractive elements exhibit some dispersion. The spectral distortion can be readily predicted and compensated given that the temperature of the refractive scanning elements and the relationship between dispersion and temperature are known.

Doyle has found that it is not necessary to perform a correction for variation of temperature when using KBr as a refractive scanning element. Nevertheless, there are materials such as germanium for which it would be necessary. It is thought that for thermal and stability, the interferometer should be enclosed in a rugged metal housing. If the outside of the housing is even lightly insulated, the high thermal conductivity of the metal will tend to make the inside isothermal. In the preferred embodiment, the spectrometer has a small thermocouple or resistance temperature detector built into it so that the temperature of the metal shell can be continuously monitored. The thermal mass of the metal housing will insure that whatever temperature variations do occur are relatively slow so that corrections at each point in time will be accurate.

The form of the error, in the spectrum, caused by dispersion is distortion of the x-axis, i.e., the wavenumber axis. If there were no dispersion, this axis would be as accurate as the reference laser stability. In order to adjust the spacing of the x-axis, it is necessary to fit a polynomial or spline to the data. The polynomial is then evaluated at unevenly spaced points along the x-axis which correspond to a grid of evenly spaced points in the absence of dispersion. This correction may be repeated for each temperature. The dispersion curve can be determined from literature coefficients, or from a NIST-traceable wavenumber absorption standard. The traceable standard is preferred.

Figure 6 diagrams the distortion of the axis and shows schematically the appropriate signal processing steps to remove the distortion. Variations of the peaks on the spectrum are grossly exaggerated for clarity. The remapping function can be readily calculated from the dispersion curve for the refractive scanning material at the operating temperature. The laser data give a very accurate indication of the thickness of the material. The modulation frequency in the top trace of Figure 6 is on a linear scale and is referenced. As the refractive index of the material increases, the modulation frequency increases in exact proportion. To produce the bottom trace on Figure 6 it is only necessary to modify the x-axis of the top trace by applying the remapping function. The remapping function simply adjusts the value of each frequency point by dividing by the refractive index at that point to produce an accurate wavenumber value. This leaves the x-axis data points unevenly spaced. If an evenly spaced wavenumber axis is desired, as is the convention in chemometric analysis, the next step is to interpolate to find an evenly spaced set of values.

Signal Processing to Record and Invert Transfer Functions

To effect the type of signal processing described above most effectively, the transfer functions of the sample and reference (generally, infrared and laser) channels must be recorded and inverted. A previous disclosure describes the use of LED's to probe the transfer functions.

Figure 7 shows a powerful method of recording and simultaneously inverting the transfer functions of the sample and reference channels is to use two adaptive filters. The adaptive filters described here may be of the finite-impulse response type describe by Widrow and Stearns (ISBN 0-13-004029-0). Each weight may be updated for each iteration of the filter computation cycle. The adaptive filter outputs are compared to delayed copies of the input signals driving the LEDs (assuming that the LEDs are operated in a region of linear response to their input signals). The error signal for each iteration may be computed from the difference between the delayed version of the test signal and the output of the filter. The error is applied as feedback to update the filter weight vector by a small amount at each iteration. Thus, each adaptive filter will converge to the inverse transfer function of the input channel, including the constant delay which is the same for both channels. The spectrometer should be controlled during the convergence time so that no spurious signals are introduced into the adaptive filter. Once the filter is fully converged, it thought to be sufficient to provide brief periods for updating the filter weight vectors to account for slow variation over time. These brief periods may occur between scans.

Figure 7 details the approach. An LED or other emitter excites the detector. In a typical spectrometer, there would be both a reference laser channel and an infrared signal channel. Thus, the signal processing detailed in G would typically appear twice in an actual system. A waveform is applied to drive an LED. The output of the LED shines on the detector. The waveform may be calculated, as from a pseudorandom algorithm, or may be generated from a lookup table. A copy of the waveform that is applied to the LED is delayed by an amount sufficient that the corresponding output from the detector has reached the midpoint of the filter delay vector. The

output of the filter is compared to the waveform that is applied to the LED (or a linearized version of the waveform applied to the LED). The reason that a linearized version may be needed is that the LED may have some non-linear response that can be precompensated by adjusting the input voltage/current to produce the desired output intensity waveform. In such a case, the input waveform differs slightly from the output of the LED. By comparing the output of the filter to the delayed version of the input to the detector, the adaptive feedback which corrects the filter weight vector will force the filter to converge to compensate for the detector transfer function. The input waveform from the LED must not have frequency components that lie outside the response range of the detector. If this condition is not met, the filter convergence problem will be ill-posed and the results unstable and inaccurate. This method would generally be applied between scans of an interferometer. During the time of scanning and data acquisition the filter weight vector would remain constant. Because the detector and preamp transfer functions would usually change only slowly over time, it would not be necessary to rapidly update the inverse transfer function unless adverse conditions are encountered. It is understood that the probing of the detector transfer functions would best be accomplished during times that the interferometer mirror is stopped between scans.

For probing the degree of saturation of the detector continuously during scanning, a different approach is required. Detector saturation can be readily monitored by exciting the LED of Figure 8 with a signal moderately higher in frequency than the highest frequency component expected in the signal or reference channel. In this case, the LED is exciting the detector simultaneously with the interferometric signal of interest. It is not thought that the reference channel generally requires compensation for saturation because the dynamic range of a sinusoidally modulated laser signal is quite small. In the composite signal, the higher frequency signal can then be demodulated to produce an estimate of the degree of saturation of the detector at each point in each scan. The magnitude of the recorded signal can then be modified for improved photometric accuracy by using the saturation information. To account for differences in the effect of saturation at different wavelengths multiple LED's can be used, each modulated with a different frequency so that the degree of saturation for a series of wavelengths can be simultaneously measured.

Figure 8 shows a diagram of 4 signals obtained when exciting the signal detector with both the interferometric signal and the probe signal. The top signal designated A shows an interference signal which has suffered from the effects of detector saturation. The dotted line indicates the intensity that should have been recorded. The trace B indicates the interferogram with a high frequency probe modulation present. The trace C indicates the demodulated intensity of the high frequency probe modulation that clearly dips at the location of the highest intensity in the interferogram. The degree of the dip very accurately captures the degree of saturation. Trace D represents the intensity of the interferogram after division by trace C. This is the corrected interferogram which has improved photometric accuracy.

LED's may also be used to probe the transfer functions of the elements of an array detector. Silicon CCD arrays, mercury-cadmium telluride arrays, indium antimonide arrays, and a variety of other array detectors are used with spectrometers for imaging and other applications. Typically, these detectors produce an output proportional to the integrated intensity falling on them during a time interval. Thus, the time-dependence need not be probed. However, the non-ideal intensity response characteristics of these components may be probed and compensated on an element-by-element basis. These non-ideal characteristics include non-linearity and offset. The output of the LED's or other emitters may be homogenized by passage through a diffuser, integrating sphere, light pipe, or other optical element to insure that each element of the array is equally illuminated.

The probe radiation for each element may be modulated independently as discussed below, but it is thought that use of equal illumination of the elements is a very practical method for probing non-uniformity.

An LED or other emission source is controlled by a computer processor which may be a DSP processor. The intensity can be accurately controlled in stepwise increments. By adjusting the intensity through a wide range, and recording the output of each detector element for each intensity, the gain, offset and non-linearity of each pixel can be separately calculated. When spectroscopic data are recorded, the transfer function of each element may be removed from the recorded data by correcting for gain, offset and linearity of response using the recorded transfer functions of the detector pixels. The saturation of the detector can be probed by taking interleaved frames from the detector with the LED on in one, and off in the next. This is analogous to the use of phase modulation of the interferometer in step-scan imaging.

The use of phase modulation in an infrared spectrometer may be used to reject multiplicative noise of the type associated with source intensity variation. This method requires that an image frame be recorded while the interferometer mirror is at one extreme of its motion. This signal is subtracted from the frame recorded with the mirror at the other extreme of motion. This would usually be repeated several times at each interferometer step and the results averaged. This practice will greatly reduce the effect of multiplicative signal variations of all types.

LED's may be also be used to probe detector transfer functions of photoacoustic (PA) or photothermal (PT) detectors. PA and PT detectors are intended to be very sensitive to temperature variations caused by absorption of light. Unfortunately, they are also very sensitive to environmental disturbances as well as convection and other non-ideal characteristics of their components. LED's or other solid state emitters, either singly or in arrays, may be used to probe the instantaneous transfer function of the detector and each point within the detector sensitive area. Thus, changes in responsivity caused by temperature variation, aging, density fluctuations (convection) or other causes may be compensated in real time. The principal variation needed relative to the above discussion is that to probe different points in a photothermal or photoacoustic detection system simultaneously requires a multiplexed probe signal. For example, in the case of a photoacoustic sample that is mounted within a photoacoustic cell, a suitable approach according to the present invention is to modulate a probe beam which is impinged on the full area of the detector active area.

The probe beam can be conveniently modulated by the use of a mirror array of the general type discussed by Messerschmidt (US Patent 6,034,370) and Tague (US Patent 5,923,036). Alternatively, a suitable probe beam can be generated by an array of LED elements that are assembled with hybrid circuit or integrated circuit technology. In either case, the probe beam is imaged onto the active area. Each point in the probe beam is then modulated with its own characteristic modulation pattern that may be a sinusoid of a particular frequency, or may be a modulation described by a Hadamard matrix. By demodulating the photoacoustic detector signal for each of the patterns contained in the probe beam, it would be possible to discern the local amplitude and phase response of the sample-gas-microphone coupling system. This is thought to be useful, because, particularly in step-scan Fourier-transform photoacoustic measurements, the variation of response over time can be source of spectral noise. By compensating these variations during data acquisition, much higher quality spectra can be measured in a much shorter time.

A source is collimated, then directed onto an array mirror device. The beam reflected from the array mirror device is directed into the photoacoustic sample chamber such that it illuminates the sample. By modulating each mirror pixel with a characteristic pattern, such as a sinusoid of known frequency and phase, it is possible to measure the amplitude and phase responses of the cell for each pixel independently. In the case of convection currents and hot spots that vary the temperature of the sample and the gas near it, it is thought that this information can be used to account for the variations in response. The composite signal from the detector, during a condition in which the response probing is active, but the input beam from a spectrometer is not active, is decomposed by Fourier transformation. As long as the probe signals are separated in frequency by more than the resolution of the Fourier transform, each probe amplitude and phase will be distinctly available. Ideally, all pixels are probed simultaneously so that information about the sample response is very dense. It is thought that the probing should be as quick as possible so that changes over short time scales can be compensated.

Etalon Standard for Reference Lasers

It would be very desirable to replace helium-neon lasers, which are typically used as frequency references for interferometry, with small, efficient semiconductor lasers. The power savings are up to 20 watts per laser replaced. This is a very important issue for portable interferometers. The waste heat generated by gas lasers has been determined to adversely affect the performance of some spectrometers. Unfortunately, for the most part, semiconductor lasers have a frequency stability on the order of parts per thousand, while the HeNe lasers have a frequency stability on the order of parts per million. The use of a small and inexpensive etalon (typically a quartz plate) to determine the absolute frequency of a laser is disclosed. An etalon with a 3cm^{-1} mode spacing may be useful, but smaller or larger spacings can also be used. Using an estimated refractive index of 1.5 would indicate an etalon thickness of about 111 microns. The etalon may be used with or without an array detector. The etalon may be placed anywhere in proximity to the interferometer, in the path from the laser to the laser detector.

Figure 9 shows a schematic of the optical and electronic layout appropriate for using such an etalon for fine control of the laser frequency. The source laser is controlled by a digital-to-analog converter which allows both a DC offset current to be applied, as well as an AC modulation current. The modulated output of the laser passes through an etalon and then through the interferometer. The optical output of the interferometer which has been modulated by the static position of the interferometer mirrors, as well as by passage through the etalon is then sampled by an analog-to-digital converter circuit. The resulting digitized signal is then demodulated by the DSP core which controls the overall spectrometer. The DSP calculates any corrections to the laser DC current to maintain a constant wavelength for the reference laser. The DSP also records the laser wavelength to use for correcting recorded interferometric data.

Choosing the mode spacing of the etalon is a compromise between two conflicting criteria - the first being that the mode spacing be narrow enough that there would be significant intensity modulation when the wavelength is varied by current modulation, and the other being that the modes are sufficiently separated that there is no ambiguity as to which mode the laser begins in at startup. It is possible to use an independent reference to ascertain in which mode of the etalon the laser is operating. For example, the rotational lines of water are known to great precision and are often present in mid-infrared spectra. By calculating the position of the water lines, the laser wavelength can be very accurately determined.

For semiconductor lasers, one powerful approach is to modulate the current in the laser to cause the wavelength to shift, and then decode the resulting intensity modulation to determine precisely the wavelength. A first order correction must first be applied to the data recorded from the output of the laser photodiode or photodiodes. The reason for the correction is that the intensity of the laser will vary approximately linearly with applied current. The intensity of the laser may be monitored with a second photodiode so that the laser intensity can be completely divided out of the values determined from the photodiodes.

Many modulation approaches are possible - the modulation could be slow relative to the scanning modulation and done continuously during operation, or it could be done only between scans with a fairly high frequency. As the semiconductor laser current is modulated, the shifting wavelength will lead to an intensity modulation as the laser moves over one or more etalon modes. Different types of lasers will have different tuning ranges. If the current tuning can be done only over a small wavelength range ($\ll 1$ mode) then it is still possible to extract from the ratio of the fundamental and the first harmonic the exact location of the laser wavelength on the etalon modes. If the current tuning range is relatively large (on the order of 1 mode) then it is possible to drive it to an even or odd symmetry point on one of the modes and maintain it there continuously.

The waveforms that are produced by wavelength modulation are Bessel functions. The frequency components of the various Bessel functions are very sensitive indications of the exact point in the etalon mode that the laser frequency is positioned. The amplitude of the Bessel functions is also a sensitive gauge of the wavelength range of the tuning. By comparing the amplitudes and phases of the various harmonics it is possible to determine the laser wavelength with excellent precision. This approach is capable of holding or determining the laser position to at least 1/1000 of a fringe, if not 1/3000 of a fringe or better. This approach gives a solid-state laser comparable frequency stability performance to a HeNe. The only cost incurred for each additional unit is the etalon. In the spectrometer design preferred by the author, the semiconductor laser is driven by a digital-to-analog channel from the DSP core, so no additional components would be required in the design. In the preferred design, the same DSP core is connected to the photodiode which detects the laser intensity as modulated by the interferometer. The additional frequency components which are produced by the various modulation schemes aimed at determining the laser frequency can be readily demodulated by the same DSP core. Since this DSP core, in the preferred design, also handle the interferometric data from the signal channel; any necessary corrections to the data can be readily recorded.

Figure 10 shows the waveforms that may be expected at different points in the etalon modes. For example, there are four points of pure symmetry within one cycle of the mode spacing. These points are the positive and negative extrema (A and C, respectively), which have purely even symmetry. That is, the etalon modes viewed from this point are symmetric. The two points of purely odd symmetry are the zero crossings, B and D.

Two important approaches which are thought to be complementary to the above description may also be used with lasers which cannot be conveniently wavelength modulated by current modulation. These are piezoelectric modulation of the etalon thickness, which will produce a similar effect. Another effective method to provide modulation for detection of wavelength is to use an electro-optic modulator at the output of the laser. It is thought that a standard HeNe which drifts by approximately a few ppm can be further improved for much less expense than the cost of a stabilized laser by the use of either of these methods. These methods are also applicable to stabilization or determination of absolute frequency of solid-state lasers of various types, including

semiconductor lasers. With gas lasers, the initial frequency will be known to much higher precision than the semiconductor laser. Thus, the mode spacings of an etalon may be much narrower so that the frequency determination is correspondingly more accurate.

The sensitivity to misalignment (i.e., non-perpendicularity of the etalon) is quite acceptable up to 0.1 degree where it reaches about 15 ppm/degree. At 1 degree the sensitivity is 10 times larger and completely unacceptable. Thus, the etalon must be precision mounted with respect to the semiconductor laser. If the etalon is mounted in such a way that the angle can be modulated, then it is quite straightforward to modulating the etalon tilt to find a minimum in the thickness. The variation of refractive index and thickness with temperature can be readily corrected if the system temperature is probed as described herein.

Optical nulling interferometry

Optical subtraction, or nulling, of interferometric signals has been known for many years as practiced and published by Griffiths, Mattson (US Patent 4,999,010) and others. The main advantage of optical subtraction is that the dynamic range of the measured interferogram is reduced. Dynamic range is particularly important in the near-infrared spectral region where the extra brightness of blackbody sources and the higher D* of detectors can make the dynamic range of interferograms considerably larger than 1 million:1. The performance of analog-to-digital converters was an important issue for mid-infrared spectral measurements for many years. Recently the performance of ADC's, particularly those of the delta sigma type, has reached the point that it can generally be disregarded for mid-IR measurements. However, the measurements made in the near-IR spectral region using interferometers can routinely achieve optical SNR levels well above 1 million:1 up to the range of 1 billion:1. Presently no ADCs can achieve this level of performance. It is impossible using current technology to digitize some interferometric NIR optical signals without seriously degrading the signal quality.

The use of permanently-aligned, tilt-compensated four-port (two-input, two-output) interferometers to perform optical subtraction of signals is disclosed. The two inputs are also out of phase, so that if a suitable comparison source (or sample) can be found, then the difference between the inputs can be nulled. An example of a suitable standard would be a pharmaceutical tablet which is known to be good. Such a tablet would be a suitable standard for inspecting other pharmaceutical tablets for quality control reasons. By using a near infrared probe beam it would be possible to obtain the chemical signature of each tablet, thereby verifying the composition as it passed by on a conveyor-belt or was loaded into containers. Such quality control would be particularly useful in cases where more than one type of pill was produced at the production facility. For optical subtraction, the ideal sample would be a pill or tablet of known composition. When a tablet of equal composition was position in the other beam, the output signal would null. If a tablet of different composition were placed in the probe beam, a large difference signal would be observed. Other examples of suitable samples for comparison might include different regions of the ground in optical remote sensing or a piece of tissue known to be healthy in the case of medical diagnostics.

In the past, the main problem with optical subtraction interferometry has been the extreme sensitivity to alignment of the interferometer components. Several disclosures have provided interferometers which are immune to small misalignments of their components, but it is thought that no previous disclosure has addressed the issue of providing an optical subtraction or optical nulling interferometer which is immune to minor misalignment of the components. The present disclosure extends these ideas to the problem of optical subtraction by showing arrangements of

interferometric components which are intrinsically and optically tilt-compensated in much the same way as the designs disclosed by Woodruff, Turner, and others.

Electronic subtraction of the signals from two different detectors is analogous to direct optical subtraction and is a practical and useful alternative, also known for many years. Generally, the two outputs must be out of phase for either type of subtraction to be useful.

Figure 11 shows a type of four port interferometer which is constructed from a monolithic beamsplitter and monolithic beam combiner elements to achieve permanent alignment. It will be appreciated by those skilled in the art that the path of the two beams from S1 and S2 will exactly coincide at D1 and D2 if the two blocks are made from the same thickness of material – independent of the exact tilt angles of the splitter and combiner. The two input beams from S1 and S2 pass through to the two outputs which are sampled by D1 and D2. The signals from these detector elements are amplified and subtracted electronically. If a subtraction succeeds in reducing the dynamic range of the interference signal by a factor of 10, it is roughly equivalent to gaining 3 bits of ADC performance. If the dynamic range can be reduced by a factor of 100, then slight more than 6 bits of ADC performance are gained. Currently, at audio frequencies, the best ADCs have about 24 bits of performance, with a dynamic range on the order of 1 million:1 To gain one or two orders of magnitude in performance is very significant improvement, and will continue to be until the SNR performance of ADCs has reached the point that the entire useful dynamic range of interferometric optical signals can be recorded.

Figures 11, 12, 13 and 14 show diagrams of four-port interferometers which are intrinsically immune to small misalignments of the components. The outputs may be sampled by two detectors and subtracted, as above, or they may be recombined, and then sampled by a single detector. For all three of these interferometers, to maintain optical alignment, the optical path difference must be scanned without introducing any tilt or shear into the beams. This is not a trivial constraint and the ramifications will be expounded below.

An alternative to subtracting the signals from two different detectors is to record or generate a facsimile of the average interferogram and subtract it electronically from each subsequent interferogram. At any point in time, the performance of available digital-to-analog converters is generally better than the performance of available analog-to-digital converters by several bits. Thus, the ability to playback a facsimile of an interferogram is sufficiently better than the ability to digitize the interferogram that subtraction of the facsimile would represent a significant improvement over the present state-of-the-art. Preferably, the average interferogram would be very accurately measured by recording a long series of interferograms, thereby reducing by averaging the imperfections of the analog-to-digital converter. The playback of the facsimile would be triggered by the fringe count of the reference laser so that it was very accurately placed in time and retardation relative to the signal being nulled.

The interferometers in Figures 11, 12, 13, and 14 depend on parallel reflector assemblies, or a related design may be constructed from the type of assembly described relative to Figure 1. Because no tilt or shear can be introduced into either arm of such an interferometer, a preferred method of refractive path difference scanning uses the change of refractive index of a gas in cell. A gas cell may be placed in each arm, and then the pressure in one cell may be increased while the pressure in the other is decreased. It is anticipated that in the relatively near future, electro-optic scanning of interferometer path difference will become feasible. At the time of this writing, the available materials and configurations can only produce very small path differences. As this